Managing the Data Center One Rack at a Time

No Limits Software White Paper #6
By David Cole

Table of Contents

Overview

The typical data center is an ever-changing system consisting of thousands of IT and infrastructure assets with very complex relationships between the three primary resources: space, power, and cooling. There are a wide range of tools available for data center management. These tools can assist the data center manager in areas such as asset management, real-time monitoring, capacity planning, and process management. In spite of this, data center management remains a daunting task. Why is this? Is it the sheer number of assets? Is it the difficulty in trying to juggle power and cooling resources in a non-homogeneous environment in which 2 kW racks must peacefully coexist with 20 kW racks?

Simplify data center management by breaking a complex environment into more manageable pieces – the racks.

Is there a way to make data center management easier? When faced with a complex problem, it is often beneficial to break the problem down into smaller segments. One method to address data center management is to break the large, complex environment (the data center) into more manageable segments – the racks. By breaking it down to the rack level, overall data center management is simplified and the resource requirements of IT assets are more closely administered.

This white paper addresses the rationale for managing at the rack level and discusses the concept of a rack as an intelligent device rather than simply a housing unit for IT equipment.

Why Manage at the Rack Level?

In the APC White Paper "Power and Cooling Capacity Management for Data Centers" (Rasmussen, 2011), Neil Rasmussen explains the reasoning for management at the rack level, saying

> *While having power and cooling supply and demand information at the room or facility level helps, it does not provide sufficiently detailed information to answer the questions about specific IT equipment deployments. An effective and practical level at which to measure and budget power and cooling capacity is at the rack level.*

Accurate answers to questions about changes to the IT infrastructure require information at the rack level.

David Chernicoff echoes this opinion in his white paper "The Shortcut Guide™ to Data Center Energy Efficiency" (Chernicoff, 2011):

> *Information about power and cooling at the data center/room level can be useful in answering general capacity questions. However, for accurate answers to questions about changes to the IT infrastructure, it is necessary to have information at the rack level. Having detailed information about power and cooling at the rack level provides the best, and most useful, information for capacity management.*

With dramatic growth in power consumption at the rack, measuring and managing power and cooling at the rack level becomes increasingly important in the overall management of data center capacity.

How Does Rack Density Affect Capacity Planning?

Data center capacity planning has typically been done at the data center level, where management makes decisions on the total power, cooling, and space resources needed to support the IT load. This planning often relies on the concept of "power density". Unfortunately, the term power density is inconsistent (Rasmussen, Guidelines for Specification of Data Center Power Density, 2005), and includes definitions such as the following:

- Power consumption of IT equipment / Area used by racks
- Power consumption of IT equipment / Area used by racks and their clearances
- Power consumption of IT equipment / Data center space
- Power consumption of IT, power and cooling equipment / Data center space
- Power consumption / Number of racks

Managing capacity based on the average rack power can result in stranded capacity and increased risk of shutdown due to insufficient resources.

The first four definitions are measured in watts/ft^2 or watts/m^2. The final definition is measured in kW/rack. In a homogeneous data center environment, where each of the racks has the same approximate power consumption, any of these definitions can be used for capacity planning as long as the definition is consistently applied. Unfortunately, racks can vary significantly in power consumption, from a patch panel rack drawing 0 kW to blade server racks drawing 20 kW or more.

Let's look at a data center with 50 racks and a total power consumption of 300 kW. The average power consumption per rack is calculated by dividing the total power consumption by the number of racks. In this case,

Average Power/Rack = 300 kW/50 racks = 6 kW/rack

If the data center is designed to support 6 kW per rack, the following figure illustrates the issues which can arise:

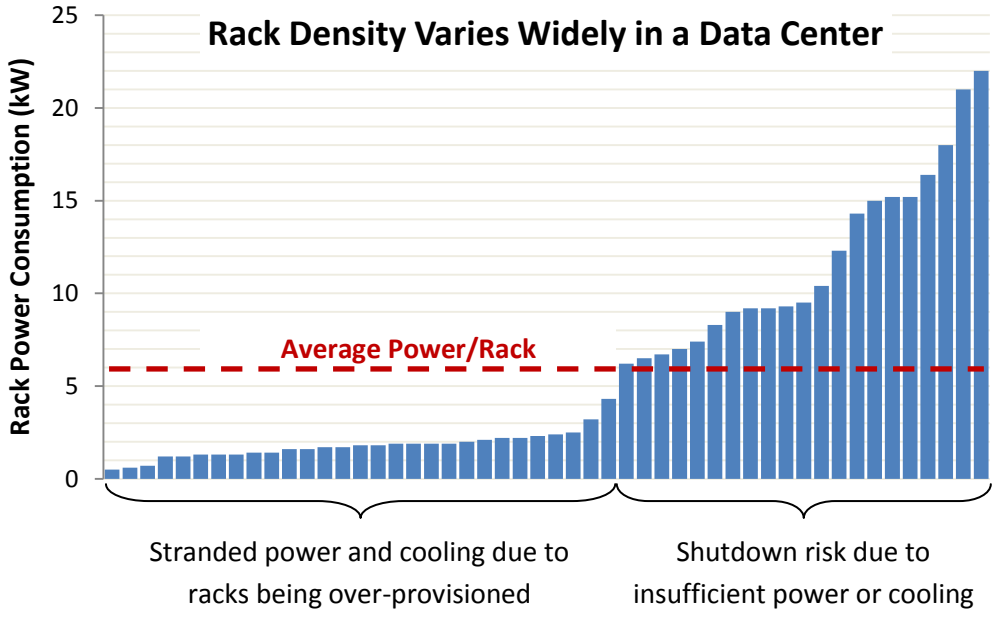

Managing to the average rack power consumption can cause two types of problems. First, there is the issue of stranded resources. If the rack is only consuming 1 kW of power but is being supplied as though it is consuming 6 kW of power, excessive power and cooling are being supplied but not used. This over-provisioning means power and cooling are being stranded and the end result is a decrease in overall energy efficiency.

It is important to understand that energy efficiency decreases for data center equipment with lower usage.

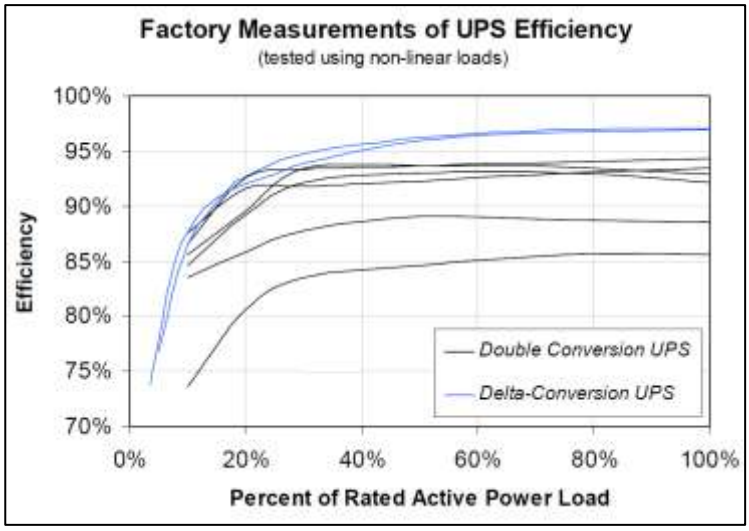

High Performance Buildings: Data Center Uninterruptible Power Supplies Berkeley Labs (Ton & Fortenbury, 2005)

This graph shows the efficiency of a UPS based on the load. A UPS loaded at 60% operates more efficiently than the same UPS loaded to only 20%. The same holds true for CRAC units, server power supplies, and so on.

Over-provisioning will always have the impact of lowering the load on the equipment and thereby reducing the overall energy efficiency.

Having sufficient power and cooling at the data center level does not mean there is sufficient power and cooling at the rack level.

The second, and potentially more serious, issue of managing to the average rack power consumption is the increased risk of shutdown due to insufficient resources. Having sufficient power and cooling at the data center level does not prevent downtime at the rack level. Power generates heat, and increasing the power consumption increases the heat generation, often creating "hot spots" within the data center. There are many ways to address hot spots including adding additional cooling, in-row cooling, segregating the data center into zones of varying densities, cold or hot aisle containment, and so on, but all require information at the rack level. Without the rack-level knowledge of power and cooling requirements, the data center manager can't make informed decisions as to where to place servers in a rack, which racks to put into a zone, which racks could benefit from cold or hot aisle containment, which rows need additional in-row cooling, and so on.

In order to properly address capacity at a rack level, data must be available at the rack level, in particular power and environmental measurements.

Measuring at the Rack Level

Once the decision is made to measure data at the rack level, the next logical question is "what should be measured". There are several key data areas to consider:

- Power
  - How much total power is the rack consuming?
  - How much power is each rack PDU providing?
  - How much power is each IT device consuming?
- Environmental Sensors
  - What is the inlet temperature at the front of the rack?
  - What is the outlet temperature at the back of the rack?
  - What is the humidity of the air being supplied to the IT equipment?
  - What is the pressure differential between the hot and cold aisles?
- Security
  - Are the front and rear rack doors open?
  - Are the front and rear rack doors locked?

"How will the information be used?" is the key question when determining what data to collect and how often it should be collected.

Not all of this data needs to be collected, of course. With most measurement considerations, the key question is "how will the information be used?" Is the data going to be used for alerting or for viewing data trends for capacity planning purposes? If the data is to be used for capacity planning purposes and the goal is to view changes over a longer period of time (say months or longer), collecting snapshot data once every hour may be sufficient. If the data is to be analyzed over a shorter period of time, say over the course of a day, snapshots should be collected more often, perhaps once every 5 minutes. If the data is to be used for alerting, however, the devices will have to be polled even more frequently or the devices will have to provide some means to "push" the data to the alerting system (via an SNMP trap, for example).

When it comes to power data, there are three primary levels of detail:

- Rack – total rack power
  - Important for knowing how much total power is consumed at the rack and how much cooling is needed to cool the rack
- Rack PDU – power at the rack PDU
  - Important for calculating the total rack power and for alerting to potential overload conditions
- Device – power consumed by the IT device
  - This information provides a greater level of detail of power at the device level but can add considerably to the cost of the rack PDU
  - This information can also be obtained through communications with the device itself using IPMI (Intelligent Platform Management

Interface) or manufacturer-specific base management protocols such as iLO or iDRAC.

Temperature readings at the rack are important to ensure that the rack inlet temperature (the temperature of the air being drawn into the IT equipment) is sufficiently low to properly cool the equipment. Just as with power, there are multiple levels of temperature monitoring which can be done at the rack level. In the case of temperature, however, the level of monitoring comes down to the number of sensors and their placement within the rack. One typical choice is to have three temperature sensors at the front of the rack – low, medium and high – in order to measure the inlet temperature variance across the front of the rack. Some choose to measure the rear temperature, particularly if using a hot-aisle containment system. Some choose to measure temperature at every other rack or only for selected racks. All of these are valid choices. More sensors will provide more temperature detail at a higher cost. Fewer sensors will be less expensive but will yield less overall information.

Typical data captured at the rack level includes power, humidity, temperature, pressure differential, and security status.

Humidity readings at the rack are important to ensure that there is sufficient moisture in the air being drawn into the IT equipment. If the humidity is too high, there is a risk of condensation forming on the equipment. If the humidity is too low, there is an increased risk of static electricity. Unlike temperature sensors, there is typically only a single humidity sensor used to measure the humidity at the front of the rack.

The pressure differential between the front and back of the rack is an important indication of the ability of the air to move from the cold aisle to the hot aisle, cooling the IT equipment in the process. If the pressure differential is too low, the rack will only be able to handle a certain load without exceeding a particular temperature threshold. This pressure differential can be measured with a pressure monitor at the front and the rear of the rack.

Security at the rack is important if only certain people are allowed to interact with equipment within the rack. The two important parameters measurements are the lock status of the rack (if employing an electronic locking mechanism) and the door open status.

Again, when determining what data is to be collected, it is important to consider how the information will be used. This will determine what sensors are needed and how often the data must be collected. Think about the key questions:

- What data do I need?
- How often do I need to measure the data?

Inventory at the Rack Level

A data center can contain thousands of assets, from servers, storage and network devices to infrastructure support equipment such as PDUs, UPSs and cooling units. Keeping track of these assets is an ongoing task faced by data center managers. A Digital Realty Trust survey found that only 26% of data center managers could locate a server that had gone down within minutes. Only 58% could locate the server within 4 hours and 20% required more than a day. The inability to locate equipment in the data center increases the mean time to repair (MTTR) for the equipment and decreases the overall availability.

20% of data center managers can take more than a day to find a server!

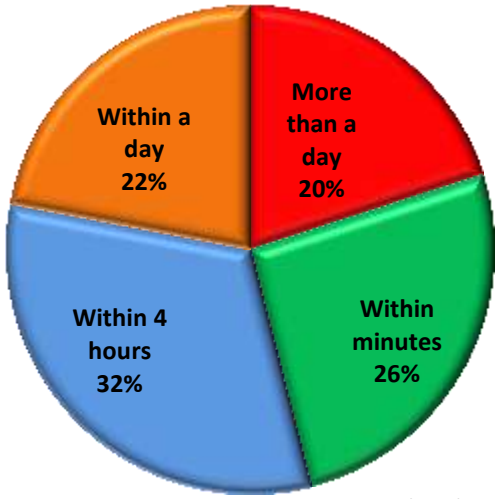

Digital Realty Trust

At the rack level, there are three primary resources which must be considered when trying to determine whether the rack can support a new asset:

- Is there enough contiguous **space** to house the asset?
- Is there sufficient redundant **power** for the asset?
- Is there enough **cooling** to remove the heat generated by the asset?

Organizations can typically expect a 10% error rate in manual data entry of asset inventory.

Comprehensive management of IT equipment must consider space, power and cooling at the rack level. Therefore, the inventory of the assets in the rack is required down to the rack unit. This is typically done through a manual data entry and audit process. There are drawbacks to this method, however. Manual entry is an arduous, time-consuming process that is typically rife with errors. In the Computer Associates technology brief *Striving to Achieve 100% Data Accuracy: The Challenge for Next Generation Asset Management* (Watson & Fulton, 2009), the authors point out the difficulty in maintaining the accuracy of this information. The authors state that "Manual tracking with pen and clipboard, or even spreadsheets is time consuming and highly error-prone. Organizations can typically expect a 10% error rate in manual data entry due to typing and transcribing errors."

There are other options for asset inventory at the rack level which can reduce or eliminate the manual process and its inherent errors. The two primary methods are radio frequency identification (RFID) or "tethered" solutions. A tethered solution ties an asset to a rack location by means of a cable or some other physical connection.

Using RFID to Track Assets within the Rack

There are multiple vendors who provide RFID-based solutions for tracking IT assets in the data center. There are some advantages to this technology, the primary one being the ability to track an asset throughout multiple locations, from the loading dock to a staging area to the rack. As the asset is moved, it is tracked using RFID readers placed throughout the building, particularly at the entrance points to a room.

RFID provides advantages over barcodes, including more storage, ability to reprogram and remote reading.

There are several major advantages of an RFID tag over a barcode. First, an RFID tag can contain much more information than a barcode, which is typically limited to 10-12 numbers or letters. Second, an RFID tag is reprogrammable while a barcode must be replaced if the information changes. Finally, an RFID tag can be read remotely while a barcode requires someone to manually scan it with a handheld device.

RFID tags are read using RFID readers which are distributed in areas in which tags are to be read. Typically, these readers are placed in doorways, enabling the readers and associated management system to determine the zone in which the tagged asset resides. As the tagged asset is moved through a doorway, the reader will detect its presence and record where the asset is currently located.

There are some drawbacks to the use of RFID in a data center. First, each tag must be programmed and installed on each asset to be tracked. Each tag is programmed with a unique identifier that is used to tie the identifier to the asset in an asset management database. This programming requires someone to gather the information about the asset and to then program this information into the tag. This manual data entry can result in errors. An alternative is to have the manufacturer program the tag with the appropriate information about the asset, including manufacturer, model and a unique asset identifier such as a serial number.

RFID is designed as a proximity system, not a location system. Tracking assets to the rack unit will require equipment to be installed in each rack.

It is important to understand that RFID was designed as a proximity or portal system rather than as a location system. In other words, the systems were designed to know when a tag passes a reader or passes through a portal containing a reader. To determine a more precise asset location, RFID must rely on triangulation of multiple readers which measure the return signal strength of a tag. This is a difficult task under any circumstances, but is particularly difficult in a data center environment with hundreds of racks and thousands of servers creating a noise-filled environment. In order to locate a server in a specific location within a rack, an RFID solution will require additional equipment to be installed in each rack. One method uses Near Field Communication, a short-range high frequency wireless communication technology which enables the exchange of data between devices over a very short distance. This method requires an RFID reader for each rack unit. A 42U rack would require 42 RFID readers, for example. These readers have a very short range (sometimes as low as 20 mm), which means there will be restrictions on where the tags can be placed on the asset in order to ensure the tag is close enough to the

reader. Proper tag mounting within the acceptable range may be difficult due to the presence of ports, power connectors and other physical limitations.

Using Tethered Solutions to Track Assets within the Rack

An alternative to RFID is to use a "tethered" solution which ties an asset to a rack location by means of a cable (typically serial or USB) or some other physical connection (one-wire, for example). These solutions are sometimes used in conjunction with an RFID solution. The combined solution uses RFID to track asset location to the room and the tethered method to get the location within the rack.

There are two types of tethered solutions. The first simply relates a unique ID to a location within the rack. A second step must then be taken to correlate the ID to a specific asset. The location system may know, for example, that asset #12345 is in rack unit 24. It won't, however, know anything about asset #12345 until there is some data entry which tells the system that the asset is an HP DL360 G6 server with serial number 9876543210.

An intelligent tethered solution can greatly reduce, or even eliminate, the manual effort required to manage asset inventory within a rack.

The second type of tethered solution is an intelligent system which communicates with the asset to retrieve detailed information about the asset and its configuration. A serial cable could be used, for example, to allow communications with a server through its serial BMC (base management controller). In a similar fashion, a USB cable could be used to allow communications with a device through a USB connection. Using this type of solution, a data center technician could place a new server in a rack and connect it via a serial or USB cable to a communications port indicating the position of the asset in the rack. Communication would be automatically established with the asset, which would then be identified. This method eliminates the manual process of linking an asset with a unique identifier. In addition, this type of solution can retrieve configuration information about the server. Detailed information such as hardware (processor, memory, disk drives, and network), operating system, serial numbers, installed software, virtual machines, and so on can be quickly and accurately retrieved from the asset. In addition, as the configuration of the server changes (software installed, hardware added, etc.), this information can be automatically logged as a change to the system.

The Rack Just Got a Lot Smarter

By breaking it down to the rack level, overall data center management can be simplified while at the same time providing the ability to more closely manage the resource requirements of IT assets. The easiest way to manage to this level is for the rack to become an intelligent device. By installing an intelligent controller and touch screen, the resultant "smart rack" can provide detailed rack-level data. With connections to existing rack management hardware, the smart rack can provide the following functionality:

- Asset management – location of all rack assets down to the rack unit
- Power management – rack PDU and overall rack power management
- Rack security – access monitoring and lock control via badge or keypad
- Environmental monitoring – temperature, humidity, air flow, and other sensors
- KVM – access to rack devices through a KVM switch
- Touch screen and keyboard at the front of the rack

Smart rack benefits include enhanced management, a single rack IP address, detailed data right at the rack, automated inventory control and the ability to easily retrofit existing racks.

Besides the obvious advantages of enhanced management of the resources needed for the IT assets, there are other key benefits to a smart rack system. First, there is a single IP address to the rack. All of the devices within the rack (rack PDU, KVM, sensors, and so on) use a private rack network for communication. This provides enhanced security while also greatly reducing the number of IP addresses required within the data center.

The touch screen on the front of the rack provides all pertinent data and functionality from a single source. Rack inventory, power usage, outlet control, environmental conditions, KVM access and rack security are all available directly from the rack. There is no need to leave the rack to gather information from a DCIM solution, spreadsheet or other tools because all of the required information is readily available. Power connections can be mapped at the rack level directly from the touch screen, so there is no need to record the connections and then later enter them into another system.

The smart rack can help to manage power by verifying the power redundancy of devices within the rack as well as warning of potential failover conditions. For example, if one rack PDU fails within the rack, the resultant failover load may cause the remaining rack PDU to also fail by exceeding its power capacity. The result is that what looked to be a redundant configuration is, in fact, dependent on a single rack PDU. The smart rack can intelligently monitor both rack PDUs and warn of potential failover conditions.

Another benefit of the smart rack is that all new assets added to the rack will be automatically added to the rack inventory (down to the rack unit) and this information can then be automatically pushed up to a centralized DCIM solution. The smart rack solution can be easily retrofitted into existing racks by communicating with equipment already installed in the rack. The controller can communicate with a wide range of components by simply extending the software capabilities.

The components of a typical smart rack system would include the following:

1) Intelligent Controller
2) Router – routes all internal network traffic on private rack network
3) Smart rack PDUs – monitor and control power to the outlet level
4) Asset Location System (both RFID or Serial solutions are supported)
    a. RFID – Near-field readers to determine asset location within the rack (requires RFID tags to be pre-configured)

b. Serial – Communicates with assets to determine asset location and other asset details (no pre-configuration required)

5) Environmental Sensors
6) Touch Screen
7) Smart Rack Interface – software which can communicate with multiple rack controllers to provide centralized control of rack security, controller upgrades, and system settings across multiple racks, as well as informing management personnel of warning and critical alerts in the data center

Summary

Smart racks can simplify data center management while also providing better management of resources at the rack level.

Even with the proper tools, managing a data center is a difficult task due to the number of assets and the complex relationships between space, power, and cooling which must be balanced. The problem can be simplified by breaking the large, complex data center environment into more manageable segments – the racks. Doing this simplifies the overall management while providing the added benefit of more closely monitoring the resource requirements of the IT equipment.

Adding intelligence to the racks allows data center managers to create "smart racks", single points of information at the critical rack level. Smart racks can provide asset management, power monitoring and control, rack security, environmental monitoring, KVM access and a single source of information at the rack itself. This information can be fed to a DCIM solution to provide an overall management solution which can address capacity planning at both data center and rack levels.

About No Limits Software

No Limits Software is a leading provider of data center infrastructure management (DCIM) solutions, including asset and change management, capacity planning and real-time monitoring.

The Rack Management Platform (RaMP) solution automates asset and change management, dramatically reduces the time to find and repair equipment, improves system availability and increases data center energy efficiency by providing accurate capacity planning.

No Limits Software is working with multiple partners to add intelligence to racks to help data center managers more effectively operate their data center.

No Limits Software was founded in 2009 by industry experts in data center monitoring and management solutions. To learn more, visit www.nolimitssoftware.com or email info@nolimitssoftware.com.

Works Cited

Chernicoff, D. (2011). *The Shortcut Guide to Data Center Energy Efficiency.* Realtime Publishers.

Rasmussen, N. (2005). *Guidelines for Specification of Data Center Power Density.* APC.

Rasmussen, N. (2011). *Power and Cooling Capacity Management for Data Centers.* APC.

Ton, M., & Fortenbury, B. (2005). *High Performance Buildings: Data Centers Uninterruptible Power Supplies (UPS).* Lawrence Berkeley National Laboratory.

Watson, T., & Fulton, J. (2009, March). *Striving to Achieve 100% Data Accuracy: The Challenge for Next Generation Asset Management.* Retrieved from http://www.ca.com/files/TechnologyBriefs/33603-itam-tech-brief-qxd-final_201662.pdf